



(19) **United States**

(12) **Patent Application Publication**
Zhang et al.

(10) **Pub. No.: US 2007/0011401 A1**

(43) **Pub. Date: Jan. 11, 2007**

(54) **SYSTEM AND METHOD FOR ADAPTIVE OPERATION OF STORAGE CAPACITIES OF RAID SYSTEMS**

(52) **U.S. Cl. 711/114**

(75) **Inventors: Ji Zhang, Monte Sereno, CA (US); Hain-Ching Liu, Fremont, CA (US); JianGang Ding, San Jose, CA (US)**

(57) **ABSTRACT**

Correspondence Address:
Irene Y.Hu
2033 Ralston Ave., #146
Belmont, CA 94002 (US)

Apparatus and methods for efficiently operating on RAID systems. A fast access buffer comprising an off-disk fast access memory module supports RAID operations such as recovery or reconfiguration operations, thereby minimizing or reducing the need for on-disk destructive zones and/or reducing disk drive I/O activities. In some cases the fast access memory module to serves as a read/write cache, reducing the need for frequent disk accesses of a small number of data blocks. Fast off-disk memory such as RAM enables rapid operation on in-buffer data blocks. Access to the material stored in the RAID devices may be enabled, partially enabled or disabled during RAID operations involving the fast access memory module and some data access operations may be synchronized with RAID operations. In some cases, data may be served from the fast access memory module, thereby providing rapid access to material stored in a RAID device during RAID operations.

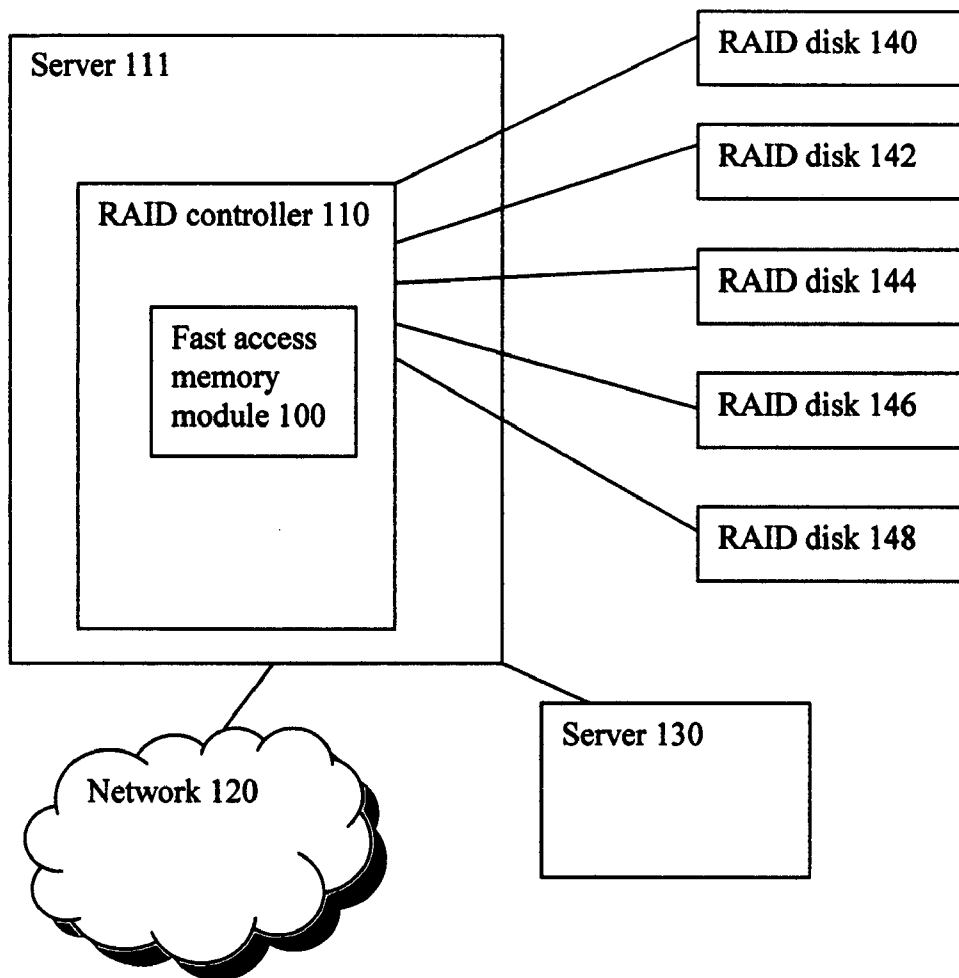
(73) **Assignee: Exavio, Inc., Santa Clara, CA**

(21) **Appl. No.: 11/175,876**

(22) **Filed: Jul. 6, 2005**

Publication Classification

(51) **Int. Cl. G06F 12/16 (2006.01)**






Figure 1

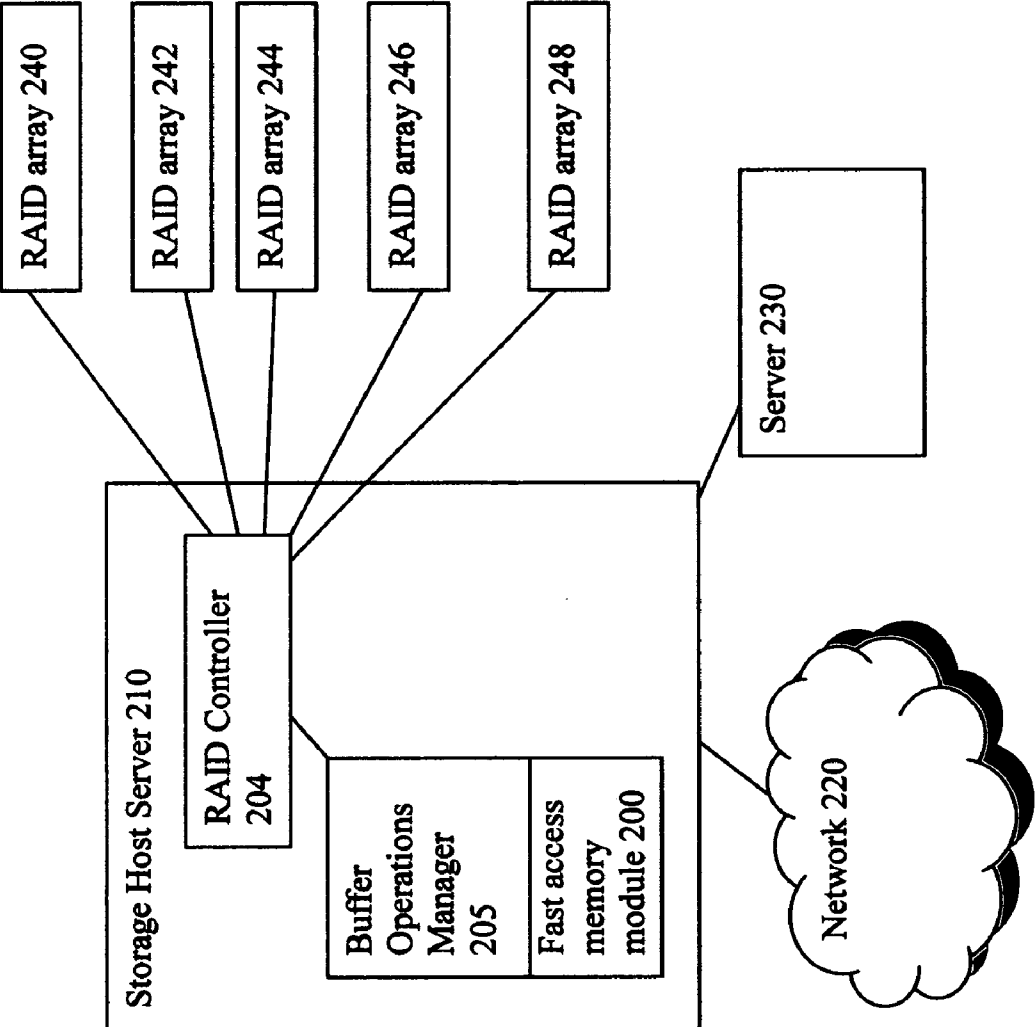


Figure 2

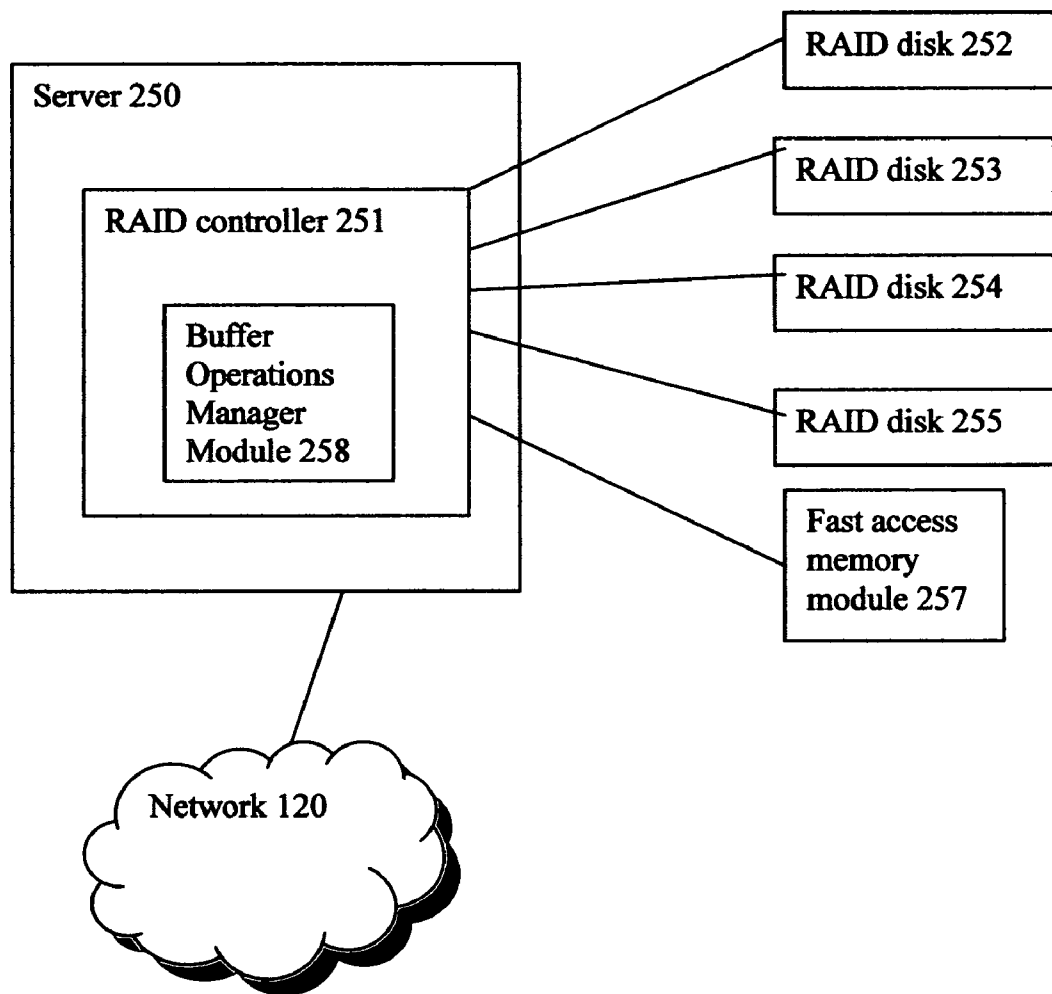


Figure 3

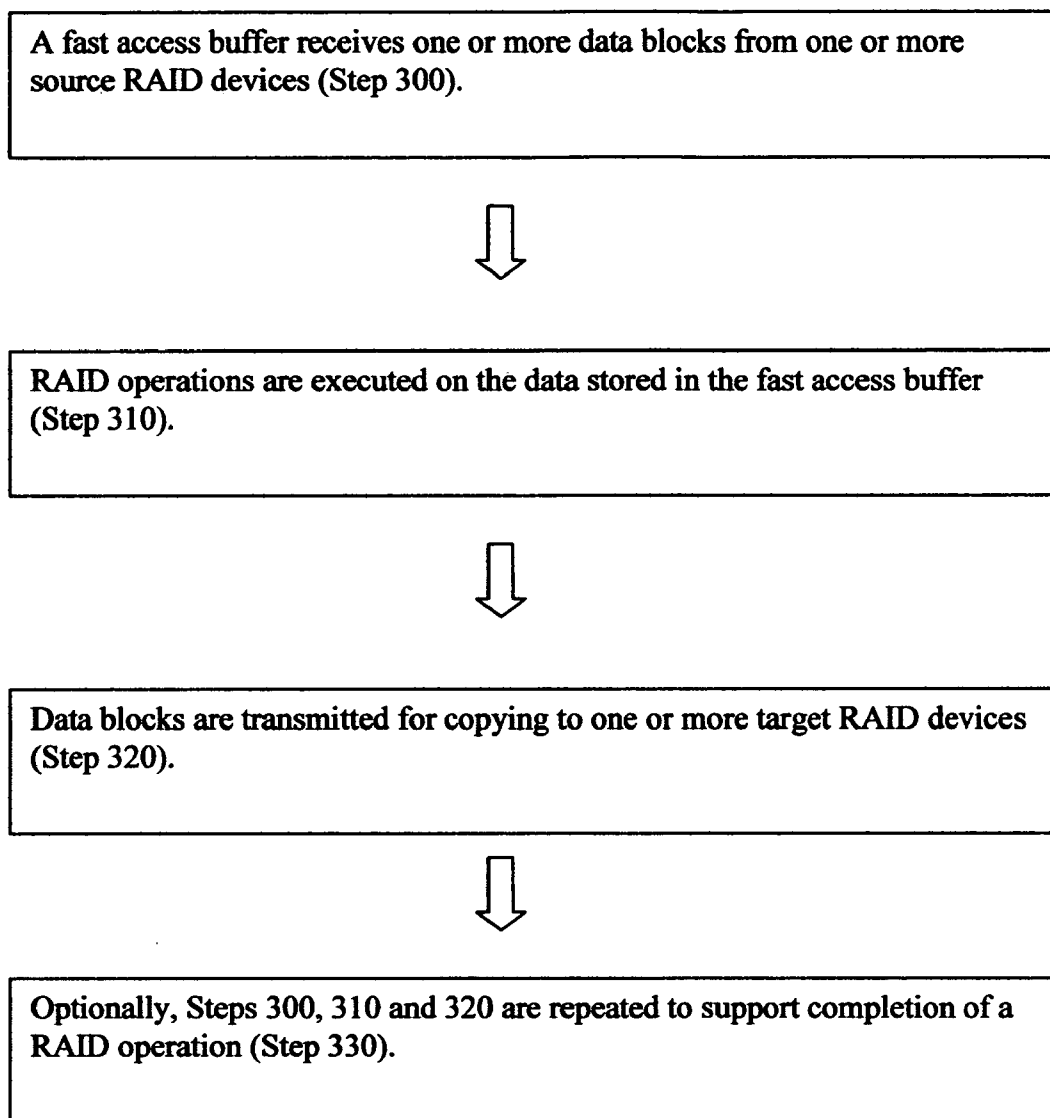


Figure 4

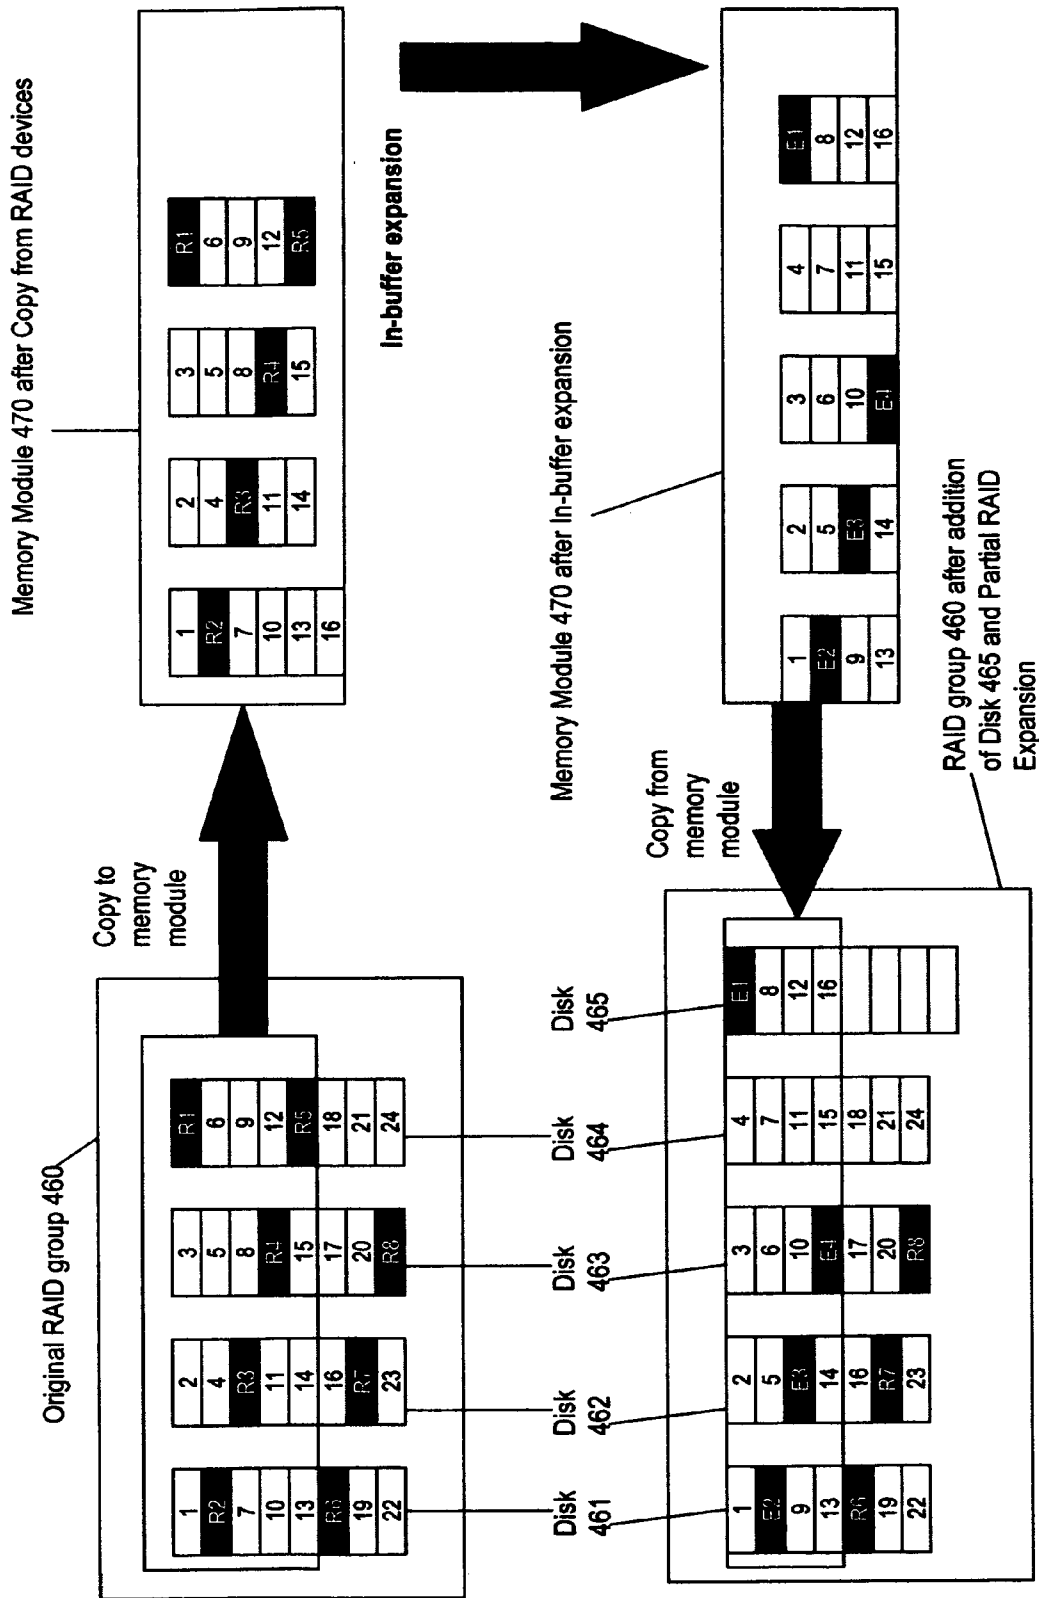


Figure 5a

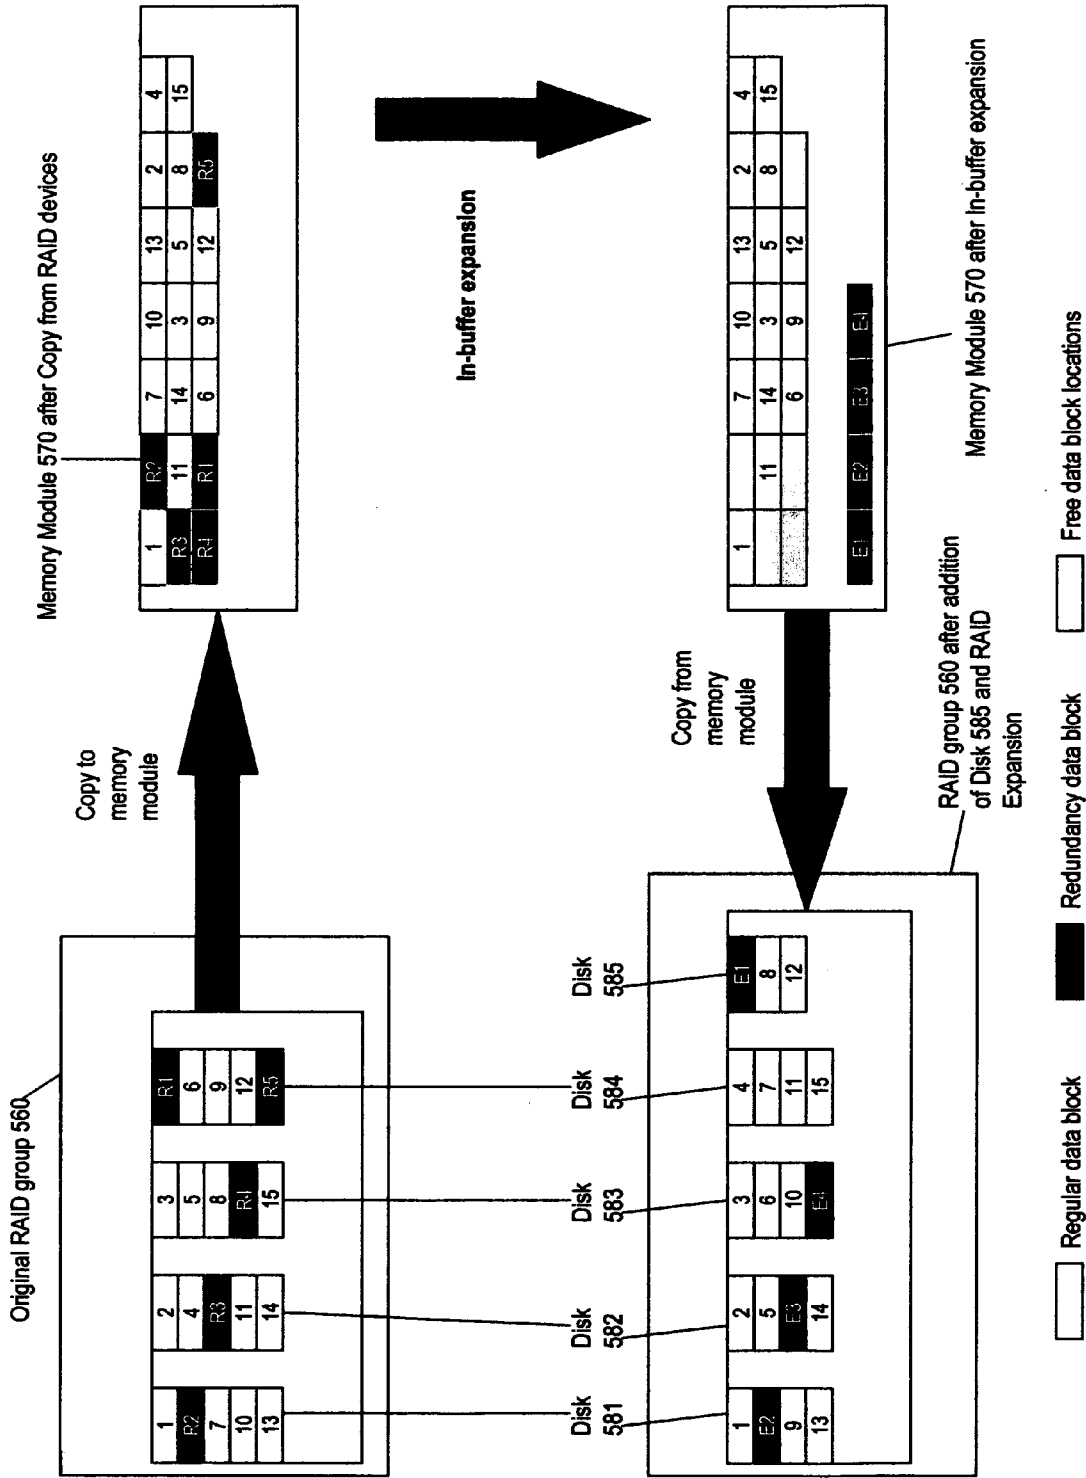


Figure 5b

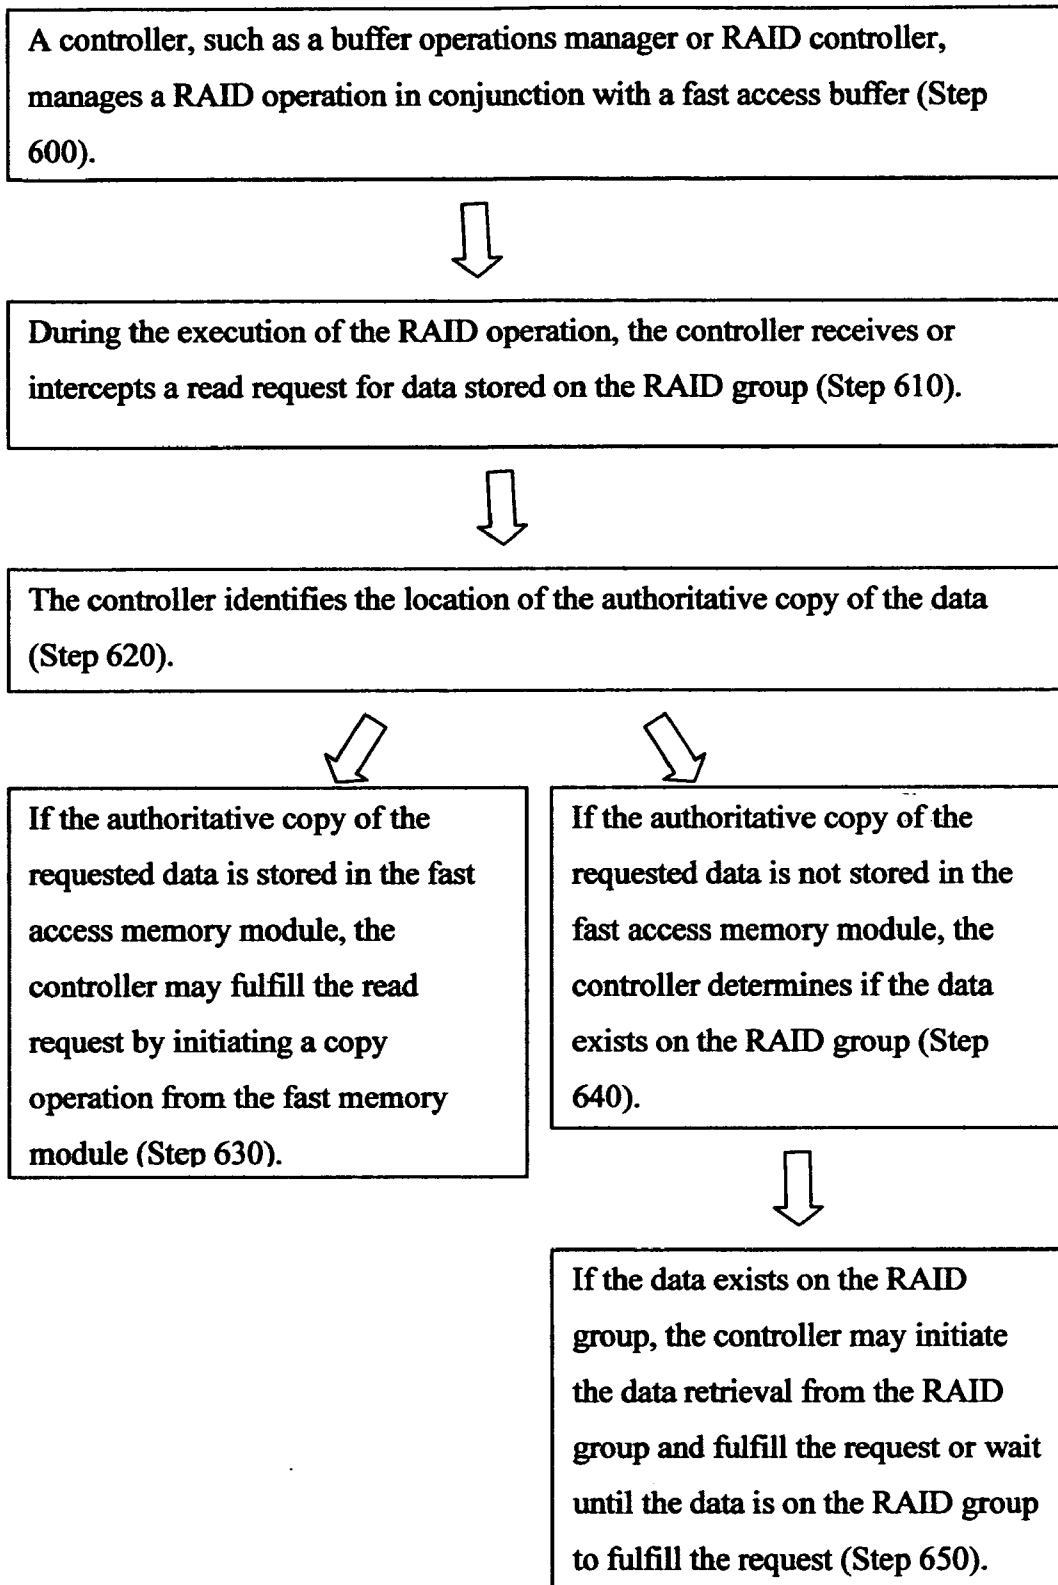


Figure 6

**SYSTEM AND METHOD FOR ADAPTIVE
OPERATION OF STORAGE CAPACITIES OF RAID
SYSTEMS**

FIELD OF INVENTION

[0001] Invention relates to apparatus and methods for performing operations on RAID (Redundant Array of Inexpensive Disks) storage systems.

BACKGROUND OF INVENTION

[0002] Storage systems using RAID (Redundant Array of Inexpensive Disks) technology are widely used. With RAID technologies, data is arranged across multiple disks (or multiple RAID arrays in a RAID-on-RAID configuration) to support redundancy and/or improved access, depending on the particular RAID configuration. When a RAID controller controls a set of RAID arrays, it is sometimes called a "RAID-on-RAID" configuration. RAID configurations require processing resources to manage the distribution of data across the disk drive storage devices or RAID arrays and, in most configurations, temporary storage or swap space to support RAID and RAID maintenance activities.

[0003] For example, in most RAID configurations, data integrity is protected by adding redundancy across multiple disk drives such that the loss of a single disk (or the loss of a RAID array in a RAID-on-RAID configuration) will not result in data loss. For example, in a RAID 5 configuration, data is stored in blocks of equal size and the data blocks are distributed across multiple disk drives. Parity data blocks are calculated and distributed across the disk drives such that if a single disk drive is lost, the contents of the lost disk drive may be rebuilt using the available data information from the other working disk drives. Rebuilding a RAID 5 disk entails deriving or calculating the contents of the lost disk drive (or RAID array in a RAID n-RAID configuration) based on the available data blocks and/or parity data blocks residing on the remaining disk drives.

[0004] In some cases, RAID operations may be very 1,0 intensive. In some cases, RAID systems may be taken off-line during RAID operations, disrupting data access. In other cases, when access to RAID systems is maintained during RAID operations, data access may be slow, due in part to the significant impact of I/O activities. For example, one I/O intensive technique for expanding a RAID 5 system includes the following steps: copying the entire data set from the existing RAID system to a back-up storage location, erasing the RAID format, adding one or more disk drive storage devices to the RAID system, reformatting the expanded group into a bigger RAID system and copying the saved data back.

[0005] In some cases the use of a destructive zone may reduce some resource requirements for expanding a RAID-5 system. For example, in some cases, a destructive zone may be enabled by allocating disk storage sections on existing disk drives in the RAID system and dedicating those sections as "destructive zones" to serve as temporary storage during a RAID expansion operations. The destructive zone(s) may be used for temporarily holding or rearranging parity data blocks, data blocks and/or intermediate calculations associated with RAID operations, thereby reducing off-disk I/O operations. However, by using an on-disk destructive zone, RAID operations may still be subject to

hard disk input/output limitations and/or hard disk access delays as data is physically relocated on a hard disk and/or copied from hard disk to hard disk. Furthermore, in some cases, access to the data stored on a RAID system may be slow or unavailable during rebuilding or maintenance activities.

[0006] In some cases, it is time consuming, disruptive and/or resource intensive to execute other RAID rebuilding or maintenance related activities such as, but not limited to, rebuilding a disk, reformatting a RAID system, expanding a RAID system, shrinking a RAID system, migrating RAID system hardware or software, changing the standard RAID system block size format.

[0007] What is needed is a method and apparatus for enabling efficient maintenance and/or recovery for RAID systems.

SUMMARY OF INVENTION

[0008] The present invention provides apparatus and methods for efficiently operating on RAID systems. For example, the current invention may be used for a variety of purposes such as, but not limited to, reconfiguring a RAID system for the purpose of modifying its storage capacity. According to the current invention, a fast access buffer comprising an off-disk fast access memory module is used to store, manipulate and/or operate on data blocks; RAID operations on the data blocks may be executed with respect to the data blocks in the fast access memory module. In some cases, multiple data blocks may be read from the source RAID device(s) and/or written to the target RAID device(s) enabling the fast access memory module to serve as a read/write cache, thereby improving input/output performance compared to RAID systems that operate based on more frequent disk accesses of a small number of data blocks.

[0009] In some cases, the fast access memory reduces or eliminates the need for on-disk storage such as, but not limited to, destructive zones or swap spaces disposed on the disk drives in a RAID configuration. The fast access memory module comprises at least some off-disk memory such as, but not limited to random access memory (RAM), and may or may not comprise additional on-disk memory. Typically, operations such as moving or manipulating data blocks in RAM take less time than analogous operations conducted in on-disk destructive zones or swap spaces established on disk drives in a RAID configuration.

[0010] In some examples according to the current invention, operation of the fast access memory module may be managed by one or more controllers such as, but not limited to, one or more controllers integrated with the fast access memory module, external controllers, server based controllers, RAID an-ay controllers and buffer operations managers. Some embodiments according to the current invention may operate in conjunction with hardware, firmware and/or software controllers or any combination thereof.

[0011] In some examples according to the current invention, during RAID operations involving the fast access memory module, access to the material stored in the RAID devices may be enabled, partially enabled or disabled. In some cases, data access operations on the RAID devices may be synchronized with RAID operations involving the

fast access memory; for example, in some cases, this may involve minimizing disk access operations, minimizing the time required to complete the RAID operation and/or minimizing delays in access operations. In some cases, when requested data is resident in the fast access memory module, it may be rapidly served from the fast access memory module, thereby providing rapid access to material stored in a RAID device during RAID operations.

[0012] In one example according to the current invention, a RAID expansion may be executed on a RAID system without requiring the data to be copied to an intermediate hard disk before copying it back to one or more disks in the expanded RAID system. In some cases, a RAID expansion may be executed without interrupting access to the data on the RAID system. Furthermore, in some cases, the I/O throughput rates for on-going read-write requests during the expansion may be maintained at relatively high rates by operating on validated copies of the data stored on the fast access memory module, thereby reducing disk access operations.

BRIEF DESCRIPTION OF DRAWINGS

[0013] FIGS. 1, 2 and 3 illustrate examples of fast access memory modules according to the current invention for enabling efficient RAID operations.

[0014] FIG. 4 is a block diagram illustrating an example process of rebuilding a single failed RAID-5 disk from a RAID 5 group using a fast access buffer according to the current invention.

[0015] FIGS. 5a and b illustrate examples of data block positions during an in-buffer expansion according to different embodiments of the current invention.

[0016] FIG. 6 is a block diagram illustrating the steps associated with an example of enabling read operations during RAID operations for an example of a fast access buffer according to the current invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[0017] According to the current invention, a fast access buffer comprises a fast access memory module and one or more interfaces for coupling the fast access memory module to one or more source RAID devices, one or more target RAID devices and one or more controllers for managing RAID operations associated with the fast access memory module. FIGS. 1, 2 and 3 illustrate examples of fast access memory modules 100, 200 and 257 according to the current invention for enabling efficient RAID operations. In the example illustrated in FIG. 1, the fast access memory module 100 comprises memory separate from disk drives 140, 142, 144, 146 and 148 such as, but not limited to, random access memory (RAM). In some cases according to the current invention, a fast access memory module may or may not comprise additional memory such as, but not limited to memory disposed on a disk drive in a RAID configuration and/or RAID device.

[0018] According to the current invention, a fast access buffer comprises one or more interfaces for coupling to one or more source RAID devices and one or more target RAID devices. A source RAID device comprises a disk drive in a RAID configuration, RAID device or RAID system which

serve as sources for the data blocks and/or redundancy data blocks received by the fast access buffer. A target RAID device comprises a disk, RAID device or RAID system which serve as destinations for the data blocks and/or redundancy data blocks output by the fast access buffer. In some cases, one or more of the source RAID device(s) may also be target RAID devices. For example, when considering the expansion of a four disk RAID array to a five disk RAID array, the four source devices would be the original four disks and the five target devices would include the original four disks in addition to the new fifth disk.

[0019] In the example illustrated in FIG. 1, fast access memory module 100 is coupled to and disposed on a RAID controller 110 which is coupled to multiple RAID devices (disk drives 140, 142, 144, 146 and 148) which may comprise source and/or target RAID devices associated with various RAID operations; in this example, fast access memory module 100 interfaces indirectly to example source RAID devices and target RAID devices (disks drives 140, 142, 144, 146 and 148) through RAID controller 110. Furthermore, in this example, RAID controller 110 controls the RAID source devices, the RAID target devices and acts as the controller for managing RAID operations associated with the fast access memory module. In this example, RAID controller 110 is implemented as a component of Server 111 and RAID devices are coupled directly to the RAID controller. However, other configurations are illustrated and envisioned.

[0020] According to the current invention, the source and/or target RAID devices may comprise disk drives in a RAID configuration and/or RAID arrays. For example, in FIG. 2, the source and target RAID devices are RAID arrays 240, 242, 244, 246 and 248 which are coupled to a fast access memory module 200 according to the current invention through Storage Host Server 210. In some cases, the source RAID devices and the target RAID devices may not be the same type of device. For example, in some cases, the current invention may be used to support a hardware migration from an array of disk drives in a RAID configuration to an array of RAID arrays wherein the source RAID devices would be disk drives in a RAID configuration and the target RAID devices would be RAID arrays.

[0021] According to various embodiments of the current invention, a fast access buffer may be used to enable efficient RAID operations related to RAID devices such as, but not limited to, rebuilding a RAID array in a RAID-on-RAID configuration, rebuilding a RAID drive, expanding a RAID array, shrinking a RAID array, migrating RAID data to accommodate a RAID hardware migration, migrating RAID data to migrate to a different RAID level, reformatting a RAID array, restoring a RAID array in a RAID-on-RAID configuration and restoring a disk drive in a RAID configuration. In some examples according to the current invention, a fast access buffer may reduce or eliminate the need for destructive zones disposed on the disk drives in a RAID configuration.

[0022] In the example illustrated in FIG. 1, the entire fast access memory module 100 is disposed on a RAID controller 110; RAID controller 110 is disposed on a server 111 and the operation of the fast access buffer is managed by the RAID controller 110. However, in other examples according

to the current invention, some or all of the fast access buffer may or may not be disposed on a RAID controller. For example, in FIG. 2, an example of a fast access memory module 200 according to the current invention is disposed on a storage host server 210; in this example, the operation of the fast access buffer 200 is managed by buffer operations manager 205.

[0023] In the example illustrated in FIG. 3, RAID controller 251 is disposed in host server 250 and fast access memory module 257 is external to both the RAID controller 251 and host server 250. In this example, the fast access memory module is managed by RAID controller 251. A buffer operations manager module 258 may be implemented in hardware, software and/or firmware. In the example illustrated in FIG. 3, the buffer operations manager module is implemented as a part of RAID controller 251.

[0024] According to alternate examples of the current invention, a buffer operations manager module may be distributed across multiple locations internal to and/or external to the RAID controller, host server and/or fast access memory module. In some cases, a buffer operations manager may be distributed across multiple software applications, hardware components, firmware components and/or hardware platforms. For example, in some cases, a buffer operations manager may use local and/or distributed server resources to execute some or all processes.

[0025] In the example illustrated in FIG. 1, fast access memory module 100 supports the operation of RAID controller 110, managing an array of disk drives 140, 142, 144, 146 and 148. According to the current invention, RAID configurations such as, but not limited to, RAID level 0, RAID level 1, RAID level 2, RAID level 3, RAID level 4, RAID level 5, RAID level 6, RAID level 10, RAID level 50, RAID level 53, RAID level 0+1 may be supported. FIG. 2 illustrates an example according to the current invention wherein a fast access buffer is used to support the operation of a RAID configuration for ensuring data integrity in case of a managed RAID array failure in a RAID-on-RAID configuration. In this example, RAID arrays 240, 242, 244, 246 and 248 each comprise their own RAID controller and two or more hard disks.

[0026] In some examples according to the current invention, a fast access buffer may reduce or eliminate the need for destructive zones disposed on the disk drives in a RAID configuration. For example, FIG. 4 is a block diagram illustrating an example process of rebuilding a single failed disk drive from a RAID-5 configured group using a fast access buffer according to the current invention. The process begins when a fast access buffer receives one or more data blocks from one or more source RAID devices (Step 300). In this example, fast access buffer may be managed by an external resource such as, but not limited to, a buffer operations manager or a RAID controller; in some cases, a fast access buffer may be integrated with a hardware, software and/or firmware management resource. In this case, a management resource initiates the retrieval of one or more data blocks from a source RAID device. In this example, the RAID operation is the rebuild of a failed RAID-5 disk which requires access to the other RAID devices in the RAID group. The management resource retrieves the data blocks from the remaining RAID devices in the RAID group, representing data and parity information, and stores the retrieved data blocks in the memory module.

[0027] In some cases, the management resource may direct the retrieval of data and parity information on a block-by-block basis. However, in some cases, the size of the memory module in the fast access buffer and the sophistication of the management resource may support copying multiple data blocks from the source RAID devices in each read operation, possibly improving performance by reducing the number of small read operations on the RAID disk drive group. In some cases, the memory module will not have the capacity to hold all of the data blocks required to complete the full RAID operation at one time and the RAID operation may be broken up into smaller RAID operations such as, but not limited to, rebuilding a segment of the failed RAID hard disk.

[0028] The process continues when RAID operations are executed on the data stored in the fast access buffer (Step 310). In this case, the RAID operation comprises regenerating missing data and parity information based on the remaining data and parity information stored on the remaining disks in the RAID disk drive group. In some cases, the fast access buffer may provide storage for intermediate calculations and/or the regenerated data and/or parity information. In some cases, supplemental storage disposed on a disk drive in a RAID configuration may or may not be used. The current invention may reduce or eliminate the need for destructive zones on disk drives in a RAID configuration; however, in some cases, the current invention may be used in conjunction with on-disk storage such as, but not limited to, destructive zones.

[0029] In some cases, the RAID operation may call for data blocks to be re-arranged on the target RAID devices. In some cases, during the execution of RAID operations on the data stored in the fast access buffer, data blocks may be re-arranged in the memory module to prepare for subsequent copying to one or more target RAID devices. However, in some cases, the target locations may be calculated, but the actual locations of the data blocks in the memory module may or may not be changed to prepare for subsequent copying to the target locations; instead, in some cases, some of the data blocks may be left in place and the subsequent copying operation may access the data blocks for copying according to their target locations, thereby copying non-adjacent data blocks in the memory module to adjacent locations on a target RAID device. When some or all of the memory module is implemented using random access memory (RAM), the order of the data blocks and/or redundancy blocks in the memory module may not be required to reflect the target order of the data blocks and/or redundancy blocks on the target RAID device. Furthermore, a random access memory (RAM) based module may support flexible and rapid operations.

[0030] The process continues when data blocks are transmitted for copying to one or more target RAID devices (Step 320). In some cases, data blocks representing data and/or parity information may be transmitted for copying to a target RAID device. In this case, the target RAID device is a replacement disk drive for the failed RAID-5 disk. In some cases, the data blocks may be copied from a storage location in the memory module of the fast access buffer. However, in other cases, the data blocks may be transmitted for copying to the target RAID device as the data blocks are calculated and they may not be associated with a temporary storage location in the fast access buffer. In some cases, copying

may take place in a block-by-block mode. However in some cases, multiple data blocks may be transmitted for copying to the target RAID device in each write operation, possibly improving performance by reducing the number of small write operations on the RAID disk drive group.

[0031] Optionally, the process continues when steps 300, 310 and/or 320 are repeated to support completion of the RAID operation (Step 330). In this example, the RAID operation is rebuilding an entire failed disk drive. Depending on the size of the memory module in the fast access buffer and the size of the failed disk drive, steps 300, 310 and 320 may be repeated multiple times.

[0032] In some embodiments of the current invention, copying data blocks to and/or from the fast access buffer may be enabled in a block-by-block method and/or a bulk method. As mentioned above, the block-by-block method may incur multiple small read and/or write operations on the RAID device(s), possibly resulting in slow performance; enabling read and/or write operations involving more than a single data block per read or write operation may reduce the total number of read and/or write operations on the RAID device(s) and may possibly result in improved performance. In some cases, enabling the largest data block size transfer supported by the RAID system and the rapid access buffer may result in the most efficient throughput.

[0033] In some cases, the fast access buffer may act as a read/write cache. In some cases, when a fast access buffer according to the current invention receives the results of a read operation involving multiple data blocks, one or more of the retrieved data blocks may be ignored for the purposes of the RAID operation; for example, data in the memory module may be used selectively wherein some data is ignored or deleted after retrieval from the source RAID device. For example, in an expansion on a RAID level 5 disk array, the original redundancy data blocks are not required to complete the expansion transaction. In some cases, the original redundancy data blocks may be retrieved, copied to the rapid access buffer and then overwritten, discarded or ignored in the rapid access buffer; in other cases, the original redundancy data blocks may not be retrieved or they may be retrieved from the hard disk, but never saved to the rapid access buffer.

[0034] As mentioned above, according to various examples of the current invention, when some or all of the memory module is implemented using random access memory (RAM), the order of the data blocks and/or redundancy blocks in the memory module may or may not reflect the target order of the data blocks and/or redundancy blocks on the target RAID device(s). In some cases, this flexibility may enable rapid RAID operations. For example, FIGS. 5a and 5b illustrate examples of data block positions during an in-buffer expansion according to different embodiments of the current invention.

[0035] FIG. 5a illustrates the position of data blocks associated with a RAID 5 level RAID group as RAID group 460 is expanded from four hard disks (460, 461, 462 and 463) to five hard disks (460, 461, 462, 463 and 464) in conjunction with a memory module 470 according to the current invention wherein the data block positions in the fast access buffer are arranged to reflect the target order of the data blocks and/or redundancy blocks on the target RAID devices. FIG. 5a is a graphic representation of the data

blocks associated with RAID group 460 and memory module 470; however, in some cases, data blocks may be aligned, but not necessarily in the same aligned format shown in FIG. 5a. In this example, the first five data blocks are read from each of the original four hard disks (460, 461, 462 and 463) and written to memory module 470. In this example, the first five data blocks comprise regular data blocks and redundancy data blocks, representing a portion of the full data set stored in RAID group 460.

[0036] The number of data blocks to read from each disk may be selected based on one or more criteria such as, but not limited to, the largest number of data blocks that may be read in a single hard disk access command, the size of the memory module and/or hard disk input/output considerations. In this example, an in-buffer expansion is executed in memory module 470, creating new redundancy blocks and discarding old redundancy blocks; data blocks are re-arranged in the memory module 470 to correspond to the target spatial position and/or target order of the data blocks on target RAID devices 460, 461, 462, 463 and 464. In this example, the arranged data blocks are copied to target hard disks 460, 461, 462, 463 and 464 based at least in part on their positions in the memory module 470; the first five data blocks from each disk represent a portion of the full data set stored in RAID group 460, so the remaining unprocessed data blocks on disks 460, 461, 462 and 463 may be processed in subsequent steps to complete the RAID system expansion.

[0037] FIG. 5b illustrates the position of data blocks associated with a RAID 5 level RAID group as RAID group 560 is expanded from four hard disks (560, 561, 562 and 563) to five hard disks (560, 561, 562, 563 and 564) in conjunction with a memory module 570 according to the current invention wherein the data block positions in the fast access buffer are not re-arranged to reflect the target order of the data blocks and/or redundancy blocks on the target RAID devices. FIG. 5b is a graphic representation of the data blocks associated with RAID group 560 and memory module 570; however, in some cases, data blocks may be aligned, but not necessarily in the same aligned format shown in FIG. 5b. In this example, the first five data blocks are read from each of the original four hard disks (560, 561, 562 and 563) and written to memory module 570. In this example, the first five data blocks comprise regular data blocks and redundancy data blocks, representing the full data set stored in RAID group 560 and the memory module 570 is large enough to hold the entire contents of the original four hard disks (560, 561, 562 and 563) as well as enough memory to support intermediate calculations and/or swap space.

[0038] The number of data blocks to read from each disk may be selected based on one or more criteria such as, but not limited to, the largest number of data blocks that may be read in a single hard disk access command, the size of the memory module and/or hard disk input/output considerations. In this example, an in-buffer expansion is executed in memory module 570, creating new redundancy blocks and discarding old redundancy blocks; data blocks are not re-arranged in the memory module 570 to correspond to the target spatial position and/or target order of the data blocks on target RAID devices 560, 561, 562, 563 and 564. This example highlights a possible benefit provided by a random access memory buffer. In this example, a controller manages

copying the data blocks from the memory module **570** to target hard disks **560**, **561**, **562**, **563** and **564**; in this example, the memory module was large enough to manage the entire contents of RAID group **560** and complete the RAID operation in a single pass.

[**0039**] In some examples of a fast access buffer according to the current invention, access to data during maintenance and/or RAID operations may be enabled. Access operations may include, but are not limited to, reading data, writing data and deleting data. For example, in some cases, the current invention may operate in a mode consistent with disabling all access, allowing partial access or allowing full access to data during the execution of maintenance and/or RAID operations. In some cases, the access mode may be managed by a RAID controller or a buffer operations manager. For example, in some cases, the current invention may be configured to disable access to data during some or all maintenance and/or RAID operations. In some cases, by disabling access to data during RAID operations, the time required to complete the RAID operation may be minimized and/or the complexity of the hardware, software and/or firmware associated with the RAID controller and/or buffer operations manager may be reduced.

[**0040**] The current invention may be used to support the operation of RAID devices, especially recovery, maintenance and reconfiguration operations. While executing RAID operations such as, but not limited to recovery, maintenance and reconfiguration operations, various embodiments of the current invention may offer various levels of concurrent access to the data on the RAID systems. Examples of three possible access levels comprise enabled, disabled or partially enabled. The access levels may or may not be the same for all operations. For example, in some cases, a system may be configured to deny all data access to a RAID system during some or all recovery, maintenance and/or reconfiguration operations; in other examples according to the current invention, partial access to data stored on the RAID systems, such as read-only access, may be enabled during some or all recovery, maintenance and/or reconfiguration operations; in other examples, full read-write access may be enabled during some or all recovery, maintenance and/or reconfiguration operations.

[**0041**] In some cases, when concurrent access is enabled, it is important to determine the location of the authoritative or master copy of data blocks. In many cases, the authoritative or master copy of data blocks is the newest copy. However, in some cases, the authoritative or master copy may not be the newest copy. For example, when data is associated with an atomic transaction, the latest copy of data blocks may not be the authoritative copy until the transaction has been committed. As RAID operations are executed, the data blocks resident on a disk drive in a RAID configuration or RAID array may not match the copy of data blocks resident in the fast access memory module. Identifying the authoritative or master copy is important in order to avoid accessing stale data and/or corrupting data stores. Furthermore, in some cases, it may be important to keep the data blocks resident in the fast access memory module in sync with the data blocks resident on the hard disks in the RAID system. For example, in some cases, the fast access memory module may be volatile. Storing the only authoritative copy of a data block in volatile memory incurs the risk of data loss or corruption in the event of power disruption. For this

reason, some examples of the current invention may require that writing data to data blocks stored in the fast access memory module be done in coordination with writing the data to data blocks stored on a disk drive managed by the RAID system.

[**0042**] In some cases, when partial or complete access to data stored on the RAID systems during RAID operations is enabled, access operations may be synchronized, prioritized and/or optimized with respect to RAID operations involving the fast access buffer and/or other activities. For example, in some cases, read-only access may be enabled during some or all RAID operations. For example, a buffer operations manager and/or RAID controller may be configured to support read operations for data stored on a RAID device and/or data temporarily stored in the fast access buffer according to the current invention.

[**0043**] In some cases, enabling access to data residing on the memory module may result in very rapid data access due in part to the fast access memory module's fast access times. For example, consider a system where read-access is enabled during a RAID-5 expansion. During the expansion, data blocks are received by the fast access buffer. If the RAID controller receives a request from a file system for data blocks that happen to be resident in the fast access buffer at the time of the receipt of that request, and if the data blocks in the fast access buffer are considered to be the authoritative copies of the data blocks, the data blocks may be provided to the file system by the fast access buffer. In some cases, this should result in very rapid data access due in part to the fast access memory module's fast access times and the relatively slower access times associated with a request for data blocks from a RAID hard disk.

[**0044**] In a similar example of synchronizing operations, consider an example according to the current invention wherein read-only access or read-write access are enabled during a RAID expansion operation. During a RAID expansion, if file system submits a request for data blocks to the RAID controller and the data blocks are not resident in the fast access buffer at the time of the receipt of the request, the data blocks are retrieved from the disk(s) drive(s) in a RAID configuration for the file system. However, the retrieved data blocks may be provided to and/or through the fast access buffer without requiring additional disk access; none, some or all of the retrieved data blocks may be used in support of processing required by the expansion operation. In this case, the scheduling of the tasks required by the RAID expansion operation may be rearranged to make timely use of the retrieved data blocks to take advantage of the retrieved data; in some cases, the usage may be selective. By synchronizing the data block retrieval for the file system with the use of the retrieved data blocks by the fast access buffer, the impact of enabling concurrent access to data while executing a RAID expansion may be reduced.

[**0045**] FIG. 6 is a block diagram illustrating the steps associated with an example of enabling read operations during RAID operations for an example of a fast access buffer according to the current invention. The process begins when a controller, such as a buffer operations manager or RAID controller, manages a RAID operation in conjunction with a fast access buffer (Step **600**). During the execution of the RAID operation, the controller receives or intercepts a read request for data stored on the RAID group (Step **610**).

The controller identifies the location of the authoritative copy of the data (Step 620). For example, the controller may check configuration parameters to confirm that when two or more copies of data exist in the environment, the copy in the fast access buffer's memory module is the authoritative copy and then check the fast access buffer's fast access memory module to determine if the requested data is temporarily stored in the fast access memory module. If the authoritative copy of the requested data is stored in the fast access memory module, the controller may fulfill the read request by initiating a copy operation from the fast access memory module (Step 630). If the authoritative copy of the requested data is not stored in the fast access memory module, the controller determines if the data exists on the RAID group (Step 640). For example, during a disk rebuild operation, the full set of data will not exist on the RAID group until the rebuild is complete. In some cases, if the data exists on the RAID group, the controller may initiate the data retrieval from the RAID group and fulfill the request or wait until the data is on the RAID group (Step 650). In some cases, coordination, synchronization and/or optimization of the RAID group access may be done to coordinate, synchronize and/or optimize access to the RAID group and/or the efficient completion of the RAID operation. For example, retrieving the data from the RAID group may be synchronized with the RAID operation to minimize access and/or copying operations associated with the disk drives in a RAID configuration.

[0046] In some cases, the current invention may be used to support adding and/or shrinking a RAID group's capacity. For example, when adding a disk to a RAID array or a new RAID array grouping a RAID-on-RAID configuration, data may be moved from disk to disk and/or RAID array to RAID array. In addition, expanding or shrinking a RAID group's capacity may require moving and/or recalculating redundancy data. For example, using the current invention to support operations such as adding and/or shrinking a RAID group's capacity may support fast access to in-buffer data, improved input/output performance based on large data transfers and reduced number of disk accesses instead of block-by-block based operations. Similarly, the current invention may be used in conjunction with RAID operations such as changing the stripe size, migrating to a different RAID level, migrating hardware and/or software.

[0047] Foregoing descriptions of specific embodiments of the invention have been presented for purposes of illustration and description. They are not intended to be exhaustive or to limit the invention to the precise forms disclosed. Modifications and variations are possible in light of the above teaching. For example, applicants contemplate that present invention may be applied for various purposes.

[0048] The embodiments were chosen and described in order to explain the principles and the application of the invention, thereby enabling others skilled in the art to utilize the invention in its various embodiments and modifications according to the particular purpose contemplated. The scope of the invention is intended to be defined by the claims appended hereto and their equivalents.

We claim:

1. A method for operating on Redundant Array of Inexpensive Disks (RAID) systems comprising the steps of:

receiving data blocks from one or more source RAID devices for storage in a fast access buffer;

executing RAID operations with respect to data blocks in the fast access buffer, and,

transmitting data blocks from the fast access buffer to one or more target RAID devices.

2. The method of claim 1 wherein:

a RAID device is selected from the list of: a disk drive, a disk array and a RAID array.

3. The method of claim 1 wherein:

the source RAID device operates according to a RAID level selected from the list of: RAID level 0, RAID level 1, RAID level 2, RAID level 3, RAID level 4, RAID level 5, RAID level 6, RAID level 10, RAID level 50, RAID level 53, RAID level 0+1; and

the target RAID device operates according to a RAID level selected from the list of: RAID level 0, RAID level 1, RAID level 2, RAID level 3, RAID level 4, RAID level 5, RAID level 6, RAID level 10, RAID level 50, RAID level 53, RAID level 0+1.

4. The method of claim 1 wherein:

at least one of the source RAID devices and at least one of the target of RAID devices are the same device.

5. The method of claim 1 wherein:

the fast access buffer comprises random access memory (RAM).

6. The method of claim 1 wherein:

the step of executing RAID operations comprises a RAID operation selected from the set of: rebuilding a RAID array, rebuilding a RAID drive, expanding a RAID array, shrinking a RAID array, migrating RAID data to accommodate a RAID hardware migration, migrating RAID data to accommodate a RAID software migration, migrating RAID data to migrate to a different RAID level, reformatting a RAID array, restoring a RAID array and restoring a disk drive in a RAID configuration.

7. The method of claim 1 wherein:

a controller is coupled to the fast access buffer; and,

wherein the step of executing RAID operations comprises the step of:

managing RAID operations in the fast access buffer, at least in part, by a controller selected from the list of: a RAID controller, a RAID controller controlling one or more source RAID devices, a RAID controller controlling one or more target RAID devices and a buffer operations manager or any combination thereof

8. The method of claim 1 wherein:

the step of executing RAID operations comprises the steps of:

identifying a data block in the fast access buffer,

calculating the target location for the identified data block on a target RAID device.

9. The method of claim 8 further comprising the steps of: after calculating the target location for the identified data block on a target RAID device, the step of copying data blocks from the fast access buffer to the one or more target RAID devices comprises the step of:

copying data blocks to the one or more target RAID devices from non-adjacent positions within the fast access buffer.

10. The method of claim 1 wherein:

the step of executing RAID operations comprises the steps of:

calculating new data redundancy information and creating one or more new data redundancy blocks reflecting the new data redundancy information; and,

identifying the target locations on the target RAID device for the new data redundancy blocks; and,

the step of copying data blocks from the fast access buffer to the one or more target RAID devices further comprises the step of writing the new data redundancy blocks to the target locations on the one or more target RAID device.

11. The method of claim 1 wherein:

the method may be executed in an off-line mode, thereby disabling data access operations during at least part of the execution of the method wherein the data access operations comprise commands initiated from an external resource targeting data stored on the RAID devices selected from the set of:

the one or more source RAID devices and the one or more target RAID devices.

12. The method of claim 1 wherein:

the method may be executed in an on-line mode whereby one or more data access operations may be successfully completed during the execution of the method wherein the data access operations comprise commands initiated from an external resource targeting RAID devices selected from the set of: the one or more source RAID devices and the one or more target RAID devices.

13. The method of claim 12 wherein:

a data access operation is an operation selected from the list of: reading data, writing data and deleting data, and,

wherein a data access operation comprises commands initiated from an external resource.

14. The method of claim 12 wherein:

the step of enabling concurrent access in an on-line mode comprises the step of:

identifying the location of the authoritative copy of the data associated with the data access operation.

15. The method of claim 12 wherein:

the data operation is writing new data to a target RAID device or deleting old data from a target RAID device; and,

the step of enabling concurrent access in an on-line mode comprises the steps of:

updating the fast access buffer according to the data operation; and,

ensuring that the target RAID device is updated according to the data operation.

16. The method of claim 12 wherein:

the step of enabling concurrent access in an on-line mode comprises the step of enabling at least one access operation to be executed on the fast access buffer for data available in the fast access buffer.

17. The method of claim 16 wherein:

the step of enabling concurrent access in an on-line mode comprises the step of synchronizing RAID operations with access operations on data on one or more RAID devices selected from the one or more source RAID devices and the one or more target RAID devices.

18. A system for efficient operation of Redundant Array of Inexpensive Disks (RAID) systems using a fast access buffer comprising:

a fast access memory module for storing data blocks; and, an interface for coupling to:

one or more source RAID devices and,

one or more target RAID devices; and,

a controller wherein the controller manages RAID operations with respect to data blocks stored in the fast access memory.

19. The apparatus of claim 18 wherein the controller is:

a RAID controller,

a RAID controller controlling a source RAID device;

a RAID controller controlling a target RAID device;

a RAID controller controlling a source RAID device and a target RAID device; or

a buffer operations manager.

20. The apparatus of claim 18 wherein the fast access memory module comprises random access memory (RAM).

* * * * *